

Identifying Reliable Annotations for Large Scale Image Segmentation

Alexander Kolesnikov
IST Austria
akolesnikov@ist.ac.at

Christoph H. Lampert
IST Austria
chl@ist.ac.at

Abstract

Challenging computer vision tasks, in particular semantic image segmentation, require large training sets of annotated images. While obtaining the actual images is often unproblematic, creating the necessary annotation is a tedious and costly process. Therefore, one often has to work with unreliable annotation sources, such as Amazon Mechanical Turk or (semi-)automatic algorithmic techniques.

In this work, we present a Gaussian process (GP) based technique for simultaneously identifying which images of a training set have unreliable annotation and learning a segmentation model in which the negative effect of these images is suppressed. Alternatively, the model can also just be used to identify the most reliably annotated images from the training set, which can then be used for training any other segmentation method.

By relying on "deep features" in combination with a linear covariance function, our GP can be learned and its hyperparameter determined efficiently using only matrix operations and gradient-based optimization. This makes our method scalable even to large datasets with several million training instances.

1. Introduction

The recent emergence of large image datasets has led to drastic progress in computer vision. In order to achieve state-of-the-art performance for various visual tasks, models are trained from millions of annotated images [14, 28]. However, manually creating expert annotation for large datasets requires a tremendous amount of resources and is often impractical, even with support by major industrial Internet companies. For example, it has been estimated that creating bounding box annotation for object detection tasks takes 25 seconds per box [27], and several minutes of human effort per image can be required to create pixel-wise annotation for semantic image segmentation tasks [16].

In order to facilitate the data annotation process and make it manageable, researchers often utilize sources of annotation that are less reliable but that scale more easily to

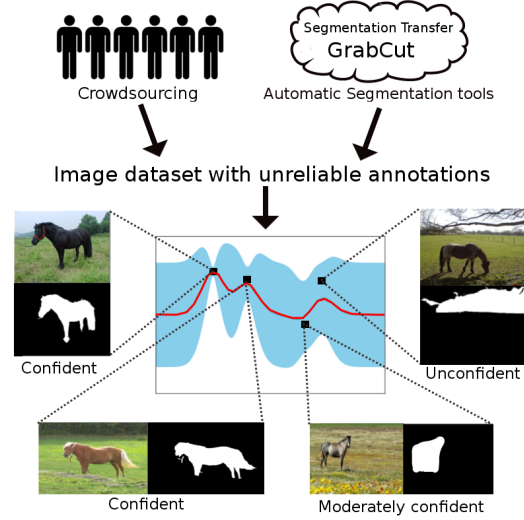


Figure 1: *Learning with unreliable image annotations*: our Gaussian process based method jointly estimates a distribution over prediction models (blue area: 95% confidence region, red line: most likely model), and a confidence value for the annotation of each image in the training data.

large amounts of data. For example, one harvests images from Internet search engines [24] or uses *Amazon Mechanical Turk* (MTurk) to create annotation. Another approach is to create annotation in a (semi-)automatic way, e.g. using knowledge transfer methods [10, 11].

A downside of such cheap data sources, in particular automatically created annotations, is that they can contain a substantial amount of mistakes. Moreover, these mistakes are often strongly correlated: for example, MTurk workers will make similar annotation errors in all images they handle, and an automatic tool, such as segmentation transfer, will work better on some classes of images than others.

Using such noisily annotated data for training can lead to suboptimal performance. As a consequence, many learning techniques try to identify and suppress the wrong or unreliable annotations in the dataset before training. However, this leads to a classical chicken-and-egg problem: one needs

a good data model to identify mislabeled parts of data and one needs reliable data to estimate a good model.

Our contribution is this work is a Gaussian processes (GP) [20] treatment of the problem of learning with unreliable annotation. It avoids the above chicken-and-egg problem by adopting a Bayesian approach, jointly learning a distribution of suitable models and confidence values for each training image (see Figure 1). Afterwards, we use the most likely such model to make predictions. All (hyper-)parameters are learned from data, so no model-selection over free parameter, such as a regularization constant or noise strength, is required.

We also describe an efficient and optionally distributed implementation of Gaussian processes with low-rank covariance matrix that scales to segmentation datasets with more than 100,000 images (16 million superpixel training instances). Conducting experiments on the task of foreground/background image segmentation with large training sets, we demonstrate that the proposed method outperforms other approaches for identifying unreliably annotated images and that this leads to improved segmentation quality.

1.1. Related work

The problem of unreliable annotation has appeared in the literature previously in different contexts.

For the task of dataset creation, it has become common practice to collect data from unreliable sources, such as MTurk, but have each sample annotated by more than one worker and combine the obtained labels, e.g. by a (weighted) majority vote [21, 26]. For segmentation tasks even this strategy can be too costly, and it is not clear how annotations could be combined. Instead, it has been suggested to have each image annotated only by a single worker, but require workers to first fulfill a grading task [16]. When using images retrieved from search engines, it has been suggested to make use of additional available information, e.g. keywords, to filter out mislabeled images [24].

For learning a classifier from unreliable data, the easiest option is to ignore the problem and rely on the fact that many discriminative learning techniques are to some extent robust against label noise. We use this strategy as one of the baselines for our experiments in Section 4, finding however that it leads to suboptimal results. Alternatively, outlier filtering based on the self-learning heuristic is popular: a prediction model is first trained on all data, then its outputs are used to identify a subset of the data consistent with the learned model. Afterwards, the model is retrained on the subset. Optionally, these steps can be repeated multiple times [4]. We use this idea as a second baseline for our experiments, showing that it improves the performance, but not as much as the method we propose.

Special variants of popular classification methods, such

as support vector machines and logistic regression, have been proposed that are more tolerant to label noise by explicitly modeling in the objective function the possibility of label changes. However, these usually result to more difficult optimization problems that need to be solved, and they can only be expected to work if certain assumptions about the noise are fulfilled, in particular that the label noise is statistically independent between different training instances. For an in-depth discussion of these and more methods we recommend the recent survey on *learning with label noise* [8].

Note that recently a method has been proposed by which a classifier is able to self-assess the quality of its predictions [30]. While also based on Gaussian processes this work differs significantly from ours: it aims at evaluating outputs of a learning system using a GP’s posterior distribution, while in this work our goal is to assess the quality of inputs for a learning system, and we do so using the GP’s ability to infer hyperparameters from data.

2. Learning with unreliable annotations

We are given a training set, $\mathcal{D} = \{(\mathbf{I}_j, \mathbf{M}_j)\}_{j=1}^n$, that consists of n pairs of images and segmentations masks. Each image \mathbf{I}_j is represented as a collection of r_j superpixels, (x_1, \dots, x_{r_j}) , with $x_k \in \mathcal{X}$ for each $k \in \{1 \dots r_j\}$, where \mathcal{X} is a universe of superpixels. Correspondingly, any segmentation mask \mathbf{M}_j is a collection (y_1, \dots, y_{r_j}) , where each $y_j \in \mathcal{Y}$ is the semantic label of the superpixel x_j and \mathcal{Y} is a finite label set. For convenience we combine all superpixels and semantic labels from the training data and form vectors \mathbf{X} and \mathbf{y} of length N , denoting individual superpixels and semantic labels by a lower index i . In the scope of this work we consider foreground-background segmentation problem with $\mathcal{Y} = \{+1, -1\}$, where $+1$ stands for foreground and -1 for background. An extension of our technique to the multiclass scenario is possible, but beyond the scope of this manuscript.

The main goal of this work is to learn a prediction function, $f : \mathcal{X} \rightarrow \mathcal{Y}$, in presence of a significant number of mistakes in the labels of the training data. We address this learning problem using Gaussian processes.

2.1. Gaussian processes

Gaussian processes (GPs) are a prominent Bayesian machine learning technique, which in particular is able to reason about noise in data and allows principled, gradient-based hyperparameter tuning. In this section we reiterate key results from the Gaussian processes literature from a practitioner’s view. For more complete discussion, see [20].

A GP is defined by a positive-definite covariance (or kernel) function, $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, that can depend on hyperparameters θ . For any test input, \bar{x} , the GP defines a Gaussian

posterior (or predictive) distribution,

$$p(\bar{y}|\bar{x}, \mathbf{X}, \mathbf{y}, \theta) = \mathcal{N}(m(\bar{x}), \sigma(\bar{x})). \quad (1)$$

The *mean function*,

$$m(\bar{x}) = \bar{\kappa}(\bar{x})^\top \mathbf{K}^{-1} \mathbf{y}, \quad (2)$$

allows us to make predictions (by taking its sign), and the *variance*,

$$\sigma(\bar{x}) = \kappa(\bar{x}, \bar{x}) - \bar{\kappa}(\bar{x})^\top \mathbf{K}^{-1} \bar{\kappa}(\bar{x}), \quad (3)$$

reflects our confidence in this prediction, where \mathbf{K} is the $N \times N$ covariance matrix of the training data with entries $\mathbf{K}_{ij} = \kappa(\mathbf{X}_i, \mathbf{X}_j)$ for $i, j \in \{1, \dots, N\}$ and $\bar{\kappa}(\bar{x}) = [\kappa(\mathbf{X}_1, \bar{x}), \dots, \kappa(\mathbf{X}_N, \bar{x})]^\top \in \mathbb{R}^N$. Note that the mean function (2) is the same as one would obtain from *kernel ridge regression* [15], which has proven effective also for classification tasks [22].

Due to their probabilistic nature, Gaussian processes can incorporate uncertainty about labels in the training set. One assumes that the label, y , of any training example is perturbed by Gaussian noise with zero mean and variance ε^2 . Different noise variances for different examples reflect the situation in which certain example labels are more trustworthy than others.

The specific form of the GP allows us to integrate out the label noise from the posterior distribution. The integral can be computed in closed form, resulting in a new posterior distribution with mean function,

$$m(\bar{x}) = \bar{\kappa}(\bar{x})^\top \mathbf{K}_\mathcal{E}^{-1} \mathbf{y}, \quad (4)$$

and variance $\sigma(\bar{x}) = \kappa(\bar{x}, \bar{x}) - \bar{\kappa}(\bar{x})^\top \mathbf{K}_\mathcal{E}^{-1} \bar{\kappa}(\bar{x})$, for an augmented covariance matrix $\mathbf{K}_\mathcal{E} = \mathbf{K} + \mathcal{E}$, where \mathcal{E} is the diagonal matrix that contains the noise variances of all training examples¹. We consider potential hyperparameters of \mathcal{E} as a part of θ .

2.2. Hyperparameter learning

A major advantage of GPs over other regression techniques is that their probabilistic interpretation offers a principled method for hyperparameter tuning based on continuous, gradient-based optimization instead of partitioning-based techniques such as cross-validation. We treat the unknown hyperparameters as random variables and study the joint probability $p(y, \theta|\mathbf{X})$ over hyperparameters and semantic labels. Employing type-II likelihood estimation (see [20], chapter 5), we obtain optimal hyperparameters, θ^* , by solving the following optimization problem,

$$\theta^* = \operatorname{argmax}_\theta \ln p(\mathbf{y}|\theta, \mathbf{X}). \quad (5)$$

¹Alternatively, we can think of $\mathbf{K}_\mathcal{E}$ as the data covariance matrix for a modified covariance function.

The expression $p(\mathbf{y}|\theta, \mathbf{X})$ in the objective (5) is known as *marginal likelihood*. Its value and gradient can be computed in closed form,

$$\ln p(\mathbf{y}|\theta, \mathbf{X}) = -\frac{1}{2}(\mathbf{y}^\top \mathbf{K}_\mathcal{E}^{-1} \mathbf{y} + \ln |\mathbf{K}_\mathcal{E}| + N \ln(2\pi)), \quad (6)$$

$$\frac{\partial \ln p(\mathbf{y}|\theta, \mathbf{X})}{\partial \theta} = \frac{1}{2} \operatorname{tr} \left((\alpha \alpha^\top - \mathbf{K}_\mathcal{E}^{-1}) \frac{\partial \mathbf{K}_\mathcal{E}}{\partial \theta} \right), \quad (7)$$

where $\alpha = \mathbf{K}_\mathcal{E}^{-1} \mathbf{y}$, θ is any entry of θ , $\frac{\partial \mathbf{K}_\mathcal{E}}{\partial \theta}$ is an element-wise partial derivative and $|\mathbf{K}_\mathcal{E}|$ denotes the determinant of $\mathbf{K}_\mathcal{E}$. If the entries of $\mathbf{K}_\mathcal{E}$ depend smoothly on θ then the maximization problem (5) is also smooth and one can apply standard gradient-based techniques, even for high-dimensional θ (i.e. many hyperparameters). While the solution is not guaranteed to be globally optimal, since (5) is not convex, the procedure has been observed to result in good estimates which are largely insensitive to the initialization [20].

2.3. A Gaussian process with groupwise confidences

Our main contribution in this work is a new approach, GP-GC, for handling unreliably annotated data in which some training examples are more trustworthy than others. Earlier GP-based approaches either assume that the noise variance is constant for all training examples, i.e. $\mathcal{E} = \lambda \mathbf{I}$ for some $\lambda > 0$, or that the noise variance is a smooth function of the inputs, $\mathcal{E} = \operatorname{diag}(g(\mathbf{X}_1), \dots, g(\mathbf{X}_N))$, where g is also a Gaussian process function [9, 12]. Neither approach is suitable for our situation: constant noise variance makes it impossible to distinguish between more and less reliable annotations. Input-dependent noise variance can reflect only errors due to image contents, which is not adequate for errors due to an unreliable annotation process. For example, in image segmentation even identically looking superpixels need not share the same noise level if they originate from different images or were annotated by different MTurk workers.

The above insight suggests to allow for arbitrary noise levels, $\mathcal{E} = \operatorname{diag}(\varepsilon_1, \dots, \varepsilon_N)$, for all training instances. However, without additional constraints this would give too much freedom in modelling the data and lead to overfitting. Therefore, we propose to take an intermediate route, based on the idea of estimating label confidence in groups. In particular, for image segmentation problems it is sufficient to model confidences for the entire image segmentation masks, instead of confidences for every individual superpixel. We obtain such a per-image confidence scores by assuming that all superpixel labels from the same image share the same confidence value, i.e. $\varepsilon_i = \varepsilon_j$ if \mathbf{X}_i and \mathbf{X}_j belong to the same image. We treat the unknown noise levels as hyperparameters and learn their value in the way described above. Since our confidence about labels is

based on the learned noise variances, we also refer to the above procedure as “learning label confidence”. We call the resulting algorithm *Gaussian Process with Groupwise Confidences*, or GP-GC.

Note that we avoid the chicken-and-egg problem mentioned in the introduction because we simultaneously obtain hyperparameters θ , in particular the noise levels $\varepsilon = [\varepsilon_1, \dots, \varepsilon_N]$, and the predictive distribution.

2.4. Instance reweighting

For unbalanced dataset, e.g. in the image segmentation case, where the background class is more frequent than the foreground, it makes sense to balance the data before training. A possible mechanism for this is to duplicate training instances of the minority class. Done naively, however, this would unnecessarily increase the computational complexity. Instead, we propose a computational shortcut that allows to incorporate duplicate instances without overhead. Let $w \in \mathbb{N}^N$ be a vector of duplicate counts, i.e. w_i is the number of copies of the training instance \mathbf{X}_i . Elementary transformations reveal that for the mean function (4), a duplication of training instances is equivalent to changing each hyperparameter ε_i to $\varepsilon_i \sqrt{w_i}$. We denote vector of hyperparameters, where ε is scaled by squared root of entries of w as θ_w . To incorporate duplicates into the marginal likelihood we also need to scale ε by the square root of vector of duplicate counts. In addition, we need to add some terms to the marginal likelihood, resulting in the following reweighted marginal likelihood,

$$\ln p_w(\mathbf{y}|\theta) \triangleq \ln p(\mathbf{y}|\theta_w) + \frac{1}{2} \sum_{i=1}^N [\ln w_i \varepsilon_i^2 - w_i \ln \varepsilon_i^2], \quad (8)$$

where “ \triangleq ” means equality up to a constant that does not depend on the hyperparameters.

Note that the above expressions are well-defined also for non-integer weights, w , which gives us not only the possibility to increase the importance of samples, but also to decrease it, if required.

3. Efficient Implementation

Gaussian processes have a reputation for being computationally demanding. Generally, their computational complexity scales cubically with the number of training instances and their memory consumption grows quadratically, because they have to store and invert the augmented data covariance matrix, $\mathbf{K}_\mathcal{E}$. Thus, standard implementations of Gaussian processes become computationally prohibitive for large-scale datasets.

Nevertheless, if the sample covariance matrix has a low-rank structure, all necessary computations can be carried out much faster by utilizing the Morrison-Sherman-Woodbury identity and the matrix determinant lemma [17,

Corollary 4.3.1]. To benefit from this, many techniques for approximating GPs by low-rank GPs have been developed using, e.g., the Nyström decomposition [32], random sub-sampling [6], k -means clustering [33], approximate kernel feature maps [19, 29], or inducing points [3, 18].

In this work we follow the general trend in computer vision and rely on an explicit feature map (obtained from a pre-trained *deep network* [5, 25]) in combination with a linear covariance function. This allows us to develop a parallel and distributed implementation of Gaussian Processes with *exact* inference, even in the large-scale regime. Formally, we use a linear covariance function, κ , of the following form,

$$\kappa(x_1, x_2) = \phi(x_1)^\top \Sigma \phi(x_2), \quad (9)$$

where $\phi : \mathcal{X} \rightarrow \mathbb{R}^k$ is a k -dimensional feature map with $k \ll N$, and $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_k^2) \in \mathbb{R}^{k \times k}$ is a diagonal matrix of feature scales. The entries of Σ are assumed to be unknown and included in the vector of hyperparameters θ . The feature map ϕ induces a feature matrix $\mathbf{F} = [\phi(\mathbf{X}_1), \dots, \phi(\mathbf{X}_N)] \in \mathbb{R}^{k \times N}$ of the training set. As a result, the augmented covariance matrix has a special structure as sum of a diagonal and a low-rank matrix,

$$\mathbf{K}_\mathcal{E} = \mathcal{E} + \mathbf{F}^\top \Sigma \mathbf{F}. \quad (10)$$

This low-rank representation allows us to store $\mathbf{K}_\mathcal{E}$ implicitly by storing matrices \mathcal{E} , Σ and \mathbf{F} , which reduces the memory requirements from $\mathcal{O}(N^2)$ to $\mathcal{O}(Nk)$. Moreover, all necessary computations for the predictive distribution (1) and marginal likelihood (6) require only $\mathcal{O}(Nk^2)$ operations instead of $\mathcal{O}(N^3)$ [32].

Computing the gradients (7) with respect to unknown hyperparameters generally imposes a computational overhead that scales linearly with the number of hyperparameters [18, 20]. For GP-GC, however, we can exploit the homogeneous structure of the hyperparameters, $\theta = [\varepsilon, \sigma]$, where $\varepsilon = [\varepsilon_1, \dots, \varepsilon_N]$ and $\sigma = [\sigma_1, \dots, \sigma_k]$ for deriving an expression for the gradient without such overhead:

$$\nabla_\varepsilon \ln p(\mathbf{y}|\theta) = \text{diag}((\alpha \alpha^\top - \mathbf{K}_\mathcal{E}^{-1}) \mathcal{E}'), \quad (11)$$

$$\nabla_\sigma \ln p(\mathbf{y}|\theta) = \text{diag}(\mathbf{F}(\alpha \alpha^\top - \mathbf{K}_\mathcal{E}^{-1}) \mathbf{F}^\top \Sigma'), \quad (12)$$

where $\alpha = \mathbf{K}_\mathcal{E}^{-1} \mathbf{y}$ and \mathcal{E}' and Σ' are diagonal matrices formed by vectors ε and σ respectively.

The computational bottleneck of low-rank Gaussian process learning is constituted by standard linear algebra routines, in particular matrix multiplication and inversion. Thus, a significant reduction in runtime can be achieved by relying on multi-threaded linear algebra libraries or even GPUs.

3.1. Distributed implementation

Despite great improvement in performance by utilizing a low-rank covariance function and parallel matrix opera-



Figure 2: Examples of training images and their segmentation masks (marked in purple) for the two datasets used. The horizontal bars reflects the quality value GP-GC assigns to the segmentation masks: the length of the bright green stripe is proportional to the number of images in the corresponding dataset that are estimated to have lower confidence than the depicted image.

tions, Gaussian processes still remain computationally challenging for truly large datasets with high-dimensional feature maps. For example, one of the datasets we use in our experiments has more than 100,000 images, 16 million superpixels and 4,113-dimensional feature representation. Storing the feature matrix alone requires more than 512 GB RAM, which is typically not available on a single workstation, but easily achievable if the representation is distributed across multiple machines.

In order to overcome memory limitations and even further improve the computational performance we developed a distributed version of low-rank Gaussian processes. It relies on the insight is that the feature matrix \mathbf{F} itself is not required for computing the prediction function (4), the marginal likelihood (6) and its gradient (7), if an *oracle* is available for answering the following four queries:

- i. compute $\mathbf{F}v$ for any $v \in \mathbb{R}^N$,
- ii. compute $\mathbf{F}^\top u$ for any $u \in \mathbb{R}^k$,
- iii. compute $\mathbf{F}D\mathbf{F}^\top$ for any diagonal $D \in \mathbb{R}^{N \times N}$,
- iv. compute $\text{diag}(\mathbf{F}^\top A\mathbf{F})$ for any $A \in \mathbb{R}^{k \times k}$.

See Appendix A for detailed explanation. On top of such an oracle we need only $\mathcal{O}(k^2 + N)$ bytes of memory and $\mathcal{O}(k^3 + N)$ operations to accomplish all necessary computations, which is orders of magnitude faster than the original requirements of $\mathcal{O}(Nk)$ bytes and $\mathcal{O}(Nk^2)$ operations.

Implementing a distributed version of the oracle is straightforward: suppose that p computational nodes are available. We then split the feature matrix $\mathbf{F} = [\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_p]$ into p roughly equally-sized parts. Each part is stored on one of the nodes in a distribute way. All oracle operations naturally decompose with respect to the parts of the feature matrix:

- i. $\mathbf{F}v = \sum_{i=1}^p \mathbf{F}_i v_i$,
- ii. $\mathbf{F}^\top u = [(\mathbf{F}_1^\top u)^\top, \dots, (\mathbf{F}_p^\top u)^\top]^\top$,
- iii. $\mathbf{F}D\mathbf{F}^\top = \sum_{i=1}^p \mathbf{F}_i D_i \mathbf{F}_i^\top$,
- iv. $\text{diag}(\mathbf{F}^\top A\mathbf{F}) = [\text{diag}(\mathbf{F}_1^\top A\mathbf{F}_1)^\top, \dots, \text{diag}(\mathbf{F}_p^\top A\mathbf{F}_p)^\top]^\top$,

where we split the vector v and the diagonal matrix D into p parts in the same fashion as we split \mathbf{F} , obtaining v_i and D_i for all $i \in \{1, \dots, p\}$. A master node takes care of distributing the objects v , u , D and A over computational nodes. Each computational node i calculates $\mathbf{F}_i v_i$,

| Method | SVM | GP | GP-GC |
|------------------------|------|------|-------------|
| <i>HDseg</i> dataset | | | |
| Horses (19,060) | 82.5 | 82.5 | 83.7 |
| Dogs (111,668) | 80.6 | 80.5 | 81.3 |
| <i>AutoSeg</i> dataset | | | |
| Horses (9,007) | 81.2 | 80.3 | 82.5 |
| Dogs (41,777) | 77.1 | 77.1 | 79.4 |
| Cats (3,006) | 73.1 | 72.4 | 73.5 |
| Sheep (5,079) | 75.6 | 75.4 | 80.0 |

Table 1: Numerical results (per-class average accuracy in %) of GP-GC and baseline approaches. The numbers in brackets indicate the number of images in the training sets. The best numbers are in bold font, see text for details on statistical significance.

$(\mathbf{F}_i^\top u)^\top$, $\mathbf{F}_i D_i \mathbf{F}_i^\top$ and $\text{diag}(\mathbf{F}_i^\top A \mathbf{F}_i)^\top$ and sends results to the master node, which collects the results of every operation and aggregates them by taking the sum for operations (i) and (iii) or the concatenation for operations (ii) and (iv). The communication between the master node and computational nodes requires sending messages of size at most $\mathcal{O}(k^2 + N)$ bytes, which is small in relation to the size of training data.

Consequently, our distributed implementation reduces the time and per-machine memory requirements by a factor of p , at the expense of minor overhead for network communication and computations on the master node.

4. Experiments

We implemented GP-GC in Python, relying the *OpenBlas* library² for linear algebra operations and L-BFGS [2] for gradient-based optimization. The code will be made publicly available.

We perform experiments on two large-scale datasets for foreground-background image segmentation, see Figure 2 for example images.

1) *HDseg* [13]³. We use the 19,060 images of horses and 111,668 images of dogs with segmentation masks created automatically by the *segmentation transfer* method [11] for training. The test images are 241 and 306 manually segmented images of horses and dogs, respectively.

2) *AutoSeg*, a new dataset that we collated from public sources and augmented with additional annotations⁴. The training images for this dataset are taken from the *ImageNet* project⁵. There are four categories: horses (9,007 images), dogs (41,777 images), cats (3,006 images), sheep (5,079 images). All training images are annotated with seg-

²<http://openblas.net>

³<http://ist.ac.at/~akolesnikov/HDseg/>

⁴We will publish the dataset, including pre-computed features.

⁵<http://www.image-net.org>

| Prediction model | SVM Classifier | | GP Classifier | |
|------------------|----------------|------------------------|------------------------|------------------|
| | SVM margin | GP-GC confidence | SVM margin | GP-GC confidence |
| | | <i>HDseg</i> dataset | <i>HDseg</i> dataset | |
| Horses | 83.8 | 84.3 | 83.5 | 84.3 |
| Dogs | 81.7 | 82.0 | 81.2 | 81.7 |
| | | <i>AutoSeg</i> dataset | <i>AutoSeg</i> dataset | |
| Horses | 82.5 | 83.2 | 82.7 | 83.9 |
| Dogs | 79.2 | 80.9 | 79.7 | 81.2 |
| Cats | 71.9 | 72.9 | 72.5 | 73.9 |
| Sheep | 80.2 | 81.7 | 81.2 | 82.9 |

Table 2: Numerical results (per-class average accuracy in %) of training an SVM or GP on filtered data. An SVM or GP classifier is trained on 25% of the most reliable images from the dataset, as selected by the SVM margin or GP-GC filtering. The best numbers are in bold font, see text for details on statistical significance.

mentation masks generated automatically by the *GrabCut* algorithm [23] from the *OpenCV* library⁶ with default parameters. We initialize *GrabCut* with bounding boxes that were also generated automatically by the *ImageNet Auto-Annotation* method [30]⁷. The test set consist of 1001 images of horses, 1521 images of dogs, 1480 images of cats and 489 images of sheep with manually created per-pixel segmentation masks that were taken from the validation part of the *MS COCO*⁸ dataset.

As evaluation metric for both datasets we use the average class accuracy [13]: we compute the percentage of correctly classified foreground pixels and the percentage of correctly classified background pixels across all images and average both values. To assess the significance of reported results, the above single number is not sufficient. Therefore, we use a closely related quantity for this purpose: we compute an average class accuracy as above separately for every image and perform a Wilcoxon signed-rank test [31] with significance level 10^{-3} .

4.1. Image Representation

We split every image into superpixels using the *SLIC* [1] method from *scikit-image*⁹ library. Each superpixel is assigned a semantic label based on the majority vote of pixel labels inside it. For each superpixel we compute appearance-based features using the *OverFeat*[25] library¹⁰. We extract a 4096-dimensional vector from the

⁶<http://opencv.org>

⁷<http://groups.inf.ed.ac.uk/calvin/proj-imagenet/page>

⁸<http://mscoco.org/>

⁹<http://scikit-image.org>

¹⁰<http://cilvr.nyu.edu/doku.php?id=software:overfeat:start>

| Method | GP-GC | Top-1% | Top-2% | Top-5% | Top-10% | Top-15% | Top-25% | Top-50% | Top-75% |
|------------------------|-------|--------|--------|--------|---------|---------|-------------|-------------|---------|
| <i>HDSeg</i> dataset | | | | | | | | | |
| Horses | 83.7 | 82.7 | 83.0 | 83.7 | 83.9 | 84.1 | 84.3 | 84.4 | 83.9 |
| Dogs | 81.3 | 80.0 | 80.3 | 80.8 | 81.2 | 81.4 | 81.7 | 81.9 | 81.6 |
| <i>AutoSeg</i> dataset | | | | | | | | | |
| Horses | 82.5 | 82.4 | 82.8 | 83.3 | 83.8 | 83.8 | 83.9 | 83.4 | 83.6 |
| Dogs | 79.4 | 77.1 | 77.7 | 78.9 | 80.1 | 80.7 | 81.2 | 80.8 | 79.9 |
| Cats | 73.5 | 57.0 | 69.3 | 72.3 | 72.6 | 73.3 | 73.9 | 74.4 | 74.3 |
| Sheep | 80.0 | 78.1 | 80.4 | 81.6 | 82.6 | 82.5 | 82.9 | 81.4 | 79.4 |

Table 3: Numerical results (per-class average accuracy in %) of training GP model from different subsets of training data. Column “Top- $\gamma\%$ ” indicates that we select $\gamma\%$ of the most reliable images according to confidences learned by GP-GC. The best numbers are in bold font.

output of the 20th layer of the pretrained model referred to as *fast* model in the library documentation. Additionally, we add features that describe the position of a superpixel in its image. For this we split each image into a 4x4 uniform grid and describe position of each superpixels by 16 values, each one specifying the ratio of pixels from the superpixel falling into the corresponding grid cell. We also add a constant (bias) feature, resulting in an overall feature map, $\phi : \mathcal{X} \rightarrow \mathbb{R}^k$, with $k = 4113$. The features within each of the three homogeneous groups (appearance, position, constant) share the same scale hyperparameter in the covariance function (9), i.e. $\sigma_i = \sigma_j$ if the feature dimensions i and j are within the same group.

4.2. Baseline approaches

We compare GP-GC against two baselines. The first baseline is also a Gaussian process, but we assume that all superpixels have the same noise variance. All hyperparameters are again learned by type-II maximum likelihood. We refer this method simply as GP. This baseline is meant to study if a selective estimation of the confidence values indeed has a positive effect on prediction performance.

As second baseline we use a linear support vector machine (SVM), relying on the *LibLinear* implementation with squared slack variables, which is known to deliver state-of-the-art optimization speed and prediction quality [7]. For training SVM models we always perform 5-fold cross-validation to determine the regularization constant $C \in \{2^{-20}, 2^{-19}, \dots, 2^{-1}\}$.

4.3. Foreground-background Segmentation

We conduct experiments on the *HDSeg* and *AutoSeg* datasets, analyzing the potential of GP-GC for two tasks: either as a dedicated method for semantic image segmentation, or as a tool for identifying reliably annotated images, which can be used afterwards, e.g., as a training set for other approaches. For all experiments we reweight training data so that foreground and background classes are bal-

anced and all instances with the same semantic label have the same weight, but the overall weight remains unchanged, i.e. $\sum_i w_i = N$. This step removes the effect of different ratios of foreground and background labels for different datasets and their subsets.

The first set of experiments compares GP-GC with the baselines, GP and SVM, on the task of foreground-background segmentation. Numeric results are presented in Table 1. They show that GP-GC achieves best results for all datasets and all semantic classes. According to a Wilcoxon signed-rank test, GP-GC’s improvement over the baselines is significant to the 10^{-3} level in all cases.

We obtain two insights from this. First, the fact that GP-GC improves over GP confirms that it is indeed beneficial to learn different confidence hyperparameters for different images. Second, the results also confirm that classification using Gaussian process regression with gradient-based hyperparameter selection yields results comparable with other state-of-the-art classifiers, such as SVMs, whose regularization parameter have to be chosen by more tedious cross-validation procedures.

In a second set of experiments we benchmark GP-GC’s ability to suppress images with unreliable annotation. For this, we apply GP-GC to the complete training set and use the learned hyperparameter values (see Figure 2 for an illustration) to form a new data set that consists only of the 25% of images that GP-GC was most confident about. We compare this approach to SVM-based filtering similar to what has been done in the computer vision literature before [4]: we train an SVM on the original dataset and form per-image confidence values by averaging the SVM margins of the contained superpixels. Afterwards we use the same construction as above, forming a new data set from the 25% of images with highest confidence scores.

We benchmark how useful the resulting data sets are by using them as training sets for either a GP (with single noise variance) or an SVM. Table 2 shows the results. By comparing the results to Table 1, one sees that both methods

for filtering out images with unreliable annotation help the segmentation accuracy. However, the improvement from filtering using GP-GC is higher than when using the data filtered by the SVM approach, regardless of the classifiers used afterwards. This indicates that GP-GC is a more reliable method for suppressing bad annotation. According to a Wilcoxon test, GP-GC’s improvement over the other method is significant to the 10^{-3} level in 11 out of 12 cases (all except *AutoSeg sheep* for the SVM classifier).

To understand this effect in more detail, we performed another experiment: we used GP-GC to create training sets of different sizes (1% to 75% of the original training sets) and trained the GP model on each of them. The results in Table 3 show that the best results are consistently obtained when using 25%–50% of the data. For example, for the largest dataset (*HDSeg dog*), the quality of the prediction model keeps increasing up to a training set of over 55,000 images (8 million superpixels). This shows that having many training images (even with unreliable annotations) is beneficial for the overall performance and that scalability is an important feature of our approach.

5. Summary

In this work we presented, GP-GC, an efficient and parameter-free method for learning from datasets with unreliable annotation, in particular for image segmentation tasks. The main idea is to use a Gaussian process to jointly model the prediction model and confidence scores of individual annotation in the training data. The confidence values are shared within groups of examples, e.g. all superpixels within an image, and can be obtained automatically using Bayesian reasoning and gradient-based hyperparameter tuning. As a consequence there are no free parameter that need to be tuned.

In experiments on two large-scale image segmentation datasets, we showed that by learning individual confidence values GP-GC is able to better cope with unreliable annotation than other classification methods. Furthermore, we showed that the estimated confidences allow us to filter out examples with unreliable annotation, thereby providing a way to create a cleaner dataset that can afterwards be used also by other learning methods.

By relying on an explicit feature map and low-rank kernel, GP-GC training is very efficient and easily implemented in a parallel or even distributed way. For example, training with 20 machines on the *HDSeg dog* segmentation dataset, which consists of over 100,000 images (16 million superpixels), takes only a few hours.

A. Reduction to the oracle

We demonstrate that having the *oracle* from Section 3.1 is sufficient to compute the mean function (2), the marginal

likelihood (6), and its gradient (7) without access to the feature matrix \mathbf{F} itself. We highlight terms that *oracle* can compute by braces with the number of the corresponding *oracle* operation.

We first apply the Sherman-Morrison-Woodbury identity and matrix determinant lemma to the matrix $\mathbf{K}_{\mathcal{E}}$:

$$\mathbf{K}_{\mathcal{E}}^{-1} = (\mathcal{E} + \mathbf{F}^{\top} \Sigma \mathbf{F})^{-1} = \mathcal{E}^{-1} - \mathcal{E}^{-1} \mathbf{F}^{\top} C^{-1} \mathbf{F} \mathcal{E}^{-1}, \quad (13)$$

$$\ln |\mathbf{K}_{\mathcal{E}}| = \ln |\mathcal{E} + \mathbf{F}^{\top} \Sigma \mathbf{F}| = \ln |\mathcal{E}| + \ln |\Sigma| + \ln |C|, \quad (14)$$

where we denote $C = \Sigma^{-1} + \underbrace{\mathbf{F} \mathcal{E}^{-1} \mathbf{F}^{\top}}_{(iii)}$.

For convenience we introduce $\tilde{\mathbf{y}} = \mathcal{E}^{-1} \mathbf{y}$. Relying on (13) we compute the following expressions:

$$\mathbf{y}^{\top} \mathbf{K}_{\mathcal{E}}^{-1} \mathbf{y} = \mathbf{y}^{\top} \tilde{\mathbf{y}} - \underbrace{(\mathbf{F} \tilde{\mathbf{y}})^{\top}}_{(i)} C^{-1} \underbrace{(\mathbf{F} \tilde{\mathbf{y}})}_{(i)}, \quad (15)$$

$$\mathbf{F} \mathbf{K}_{\mathcal{E}}^{-1} \mathbf{y} = \underbrace{\mathbf{F} \tilde{\mathbf{y}}}_{(i)} - \underbrace{(\mathbf{F} \mathcal{E}^{-1} \mathbf{F}^{\top})}_{(iii)} C^{-1} \underbrace{(\mathbf{F} \tilde{\mathbf{y}})}_{(i)}, \quad (16)$$

$$\mathbf{F} \mathbf{K}_{\mathcal{E}}^{-1} \mathbf{F}^{\top} = \underbrace{(\mathbf{F} \mathcal{E}^{-1} \mathbf{F}^{\top})}_{(iii)} (\mathbf{I} - \underbrace{C^{-1} (\mathbf{F} \mathcal{E}^{-1} \mathbf{F}^{\top})}_{(iii)}), \quad (17)$$

$$\alpha = \mathbf{K}_{\mathcal{E}}^{-1} \mathbf{y} = \tilde{\mathbf{y}} - \mathcal{E}^{-1} \underbrace{\mathbf{F}^{\top} C^{-1} \mathbf{F} \tilde{\mathbf{y}}}_{(i)}, \quad (18)$$

$$\text{diag}(\mathbf{K}_{\mathcal{E}}^{-1}) = \text{diag}(\mathcal{E}^{-1}) - \underbrace{\text{diag}(\mathbf{F}^{\top} C^{-1} \mathbf{F})}_{(iv)} \odot \text{diag}(\mathcal{E}^{-2}), \quad (19)$$

where \odot is elementwise (Hadamard) vector multiplication.

Using the above identities, we obtain the mean of the predictive distribution,

$$m(\bar{x}) = \bar{\kappa}(\bar{x})^{\top} \mathbf{K}_{\mathcal{E}}^{-1} \mathbf{y} = \phi(\bar{x})^{\top} \underbrace{\mathbf{F} \mathbf{K}_{\mathcal{E}}^{-1} \mathbf{y}}_{(16)}, \quad (20)$$

and the marginal likelihood,

$$\ln p(\mathbf{y}|\boldsymbol{\theta}) = -\frac{1}{2} \left(\underbrace{\mathbf{y}^{\top} \mathbf{K}_{\mathcal{E}}^{-1} \mathbf{y}}_{(15)} + \underbrace{\ln |\mathbf{K}_{\mathcal{E}}|}_{(14)} + N \ln(2\pi) \right). \quad (21)$$

Finally, we compute the gradient of the marginal likelihood with respect to the noise variances ε ,

$$\begin{aligned} \nabla_{\varepsilon} \ln p(\mathbf{y}|\boldsymbol{\theta}) &= \text{diag}((\alpha \alpha^{\top} - \mathbf{K}_{\mathcal{E}}^{-1}) \mathcal{E}') \\ &= (\alpha \odot \alpha - \underbrace{\text{diag}(\mathbf{K}_{\mathcal{E}}^{-1})}_{(19)}) \odot \text{diag}(\mathcal{E}') \end{aligned} \quad (22)$$

and with respect to the feature scales σ ,

$$\begin{aligned} \nabla_{\sigma} \ln p(\mathbf{y}|\boldsymbol{\theta}) &= \text{diag}(\mathbf{F}(\alpha \alpha^{\top} - \mathbf{K}_{\mathcal{E}}^{-1}) \mathbf{F}^{\top} \Sigma') \\ &= \underbrace{(\mathbf{F} \alpha)}_{(i)} \odot \underbrace{(\mathbf{F} \alpha)}_{(i)} - \underbrace{\text{diag}(\mathbf{F} \mathbf{K}_{\mathcal{E}}^{-1} \mathbf{F}^{\top})}_{(17)} \odot \text{diag}(\Sigma'). \end{aligned} \quad (23)$$

References

- [1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, 2012.
- [2] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 1995.
- [3] J. Chen, N. Cao, K. H. Low, R. Ouyang, C. K.-Y. Tan, and P. Jaillet. Parallel Gaussian process regression with low-rank covariance matrix approximations. In *Uncertainty in Artificial Intelligence (UAI)*, 2013.
- [4] T. Deselaers, B. Alexe, and V. Ferrari. Weakly supervised localization and learning with generic knowledge. *International Journal of Computer Vision (IJCV)*, 100(3):275–293, 2012.
- [5] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. DeCAF: A deep convolutional activation feature for generic visual recognition. In *International Conference on Machine Learning (ICML)*, 2014.
- [6] P. Drineas and M. W. Mahoney. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *Journal of Machine Learning Research (JMLR)*, 6:2153–2175, 2005.
- [7] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research (JMLR)*, 2008.
- [8] B. Fréney and M. Verleysen. Classification in the presence of label noise: a survey. *IEEE Transactions on Neural Networks and Learning Systems (T-NN)*, 25(5):845–869, 2014.
- [9] P. W. Goldberg, C. K. Williams, and C. M. Bishop. Regression with input-dependent noise: A Gaussian process treatment. *Conference on Neural Information Processing Systems (NIPS)*, 10, 1997.
- [10] M. Guillaumin and V. Ferrari. Large-scale knowledge transfer for object localization in ImageNet. In *Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2012.
- [11] M. Guillaumin, D. Küttel, and V. Ferrari. ImageNet auto-annotation with segmentation propagation. *International Journal of Computer Vision (IJCV)*, 2014.
- [12] K. Kersting, C. Plagemann, P. Pfaff, and W. Burgard. Most likely heteroscedastic Gaussian process regression. In *International Conference on Machine Learning (ICML)*, 2007.
- [13] A. Kolesnikov, M. Guillaumin, V. Ferrari, and C. H. Lampert. Closed-form approximate CRF training for scalable image segmentation. In *European Conference on Computer Vision (ECCV)*, 2014.
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *Conference on Neural Information Processing Systems (NIPS)*, 2012.
- [15] C. H. Lampert. Kernel methods in computer vision. *Foundations and Trends in Computer Graphics and Vision*, 4(3):193–285, 2009.
- [16] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 2014.
- [17] K. P. Murphy. *Machine learning: a probabilistic perspective*. The MIT Press, 2012.
- [18] J. Quiñero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research (JMLR)*, 2005.
- [19] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Conference on Neural Information Processing Systems (NIPS)*, 2007.
- [20] C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. The MIT Press, 2006.
- [21] V. C. Raykar, S. Yu, L. H. Zhao, G. H. Valadez, C. Florin, L. Bogoni, and L. Moy. Learning from crowds. *Journal of Machine Learning Research (JMLR)*, 11:1297–1322, 2010.
- [22] R. Rifkin, G. Yeo, and T. Poggio. Regularized least-squares classification. In *Advances in Learning Theory: Methods, Models and Applications*, chapter 7. IOS Press, 2003.
- [23] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. *ACM Transactions on Graphics (TOG)*, 23(3):309–314, 2004.
- [24] F. Schroff, A. Criminisi, and A. Zisserman. Harvesting image databases from the web. *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, 33(4):754–766, 2011.
- [25] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. OverFeat: Integrated recognition, localization and detection using convolutional networks. [arXiv:1312.6229 \[cs.CV\]](https://arxiv.org/abs/1312.6229), 2013.
- [26] A. Sorokin and D. Forsyth. Utility data annotation with Amazon Mechanical Turk. In *Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [27] H. Su, J. Deng, and L. Fei-Fei. Crowdsourcing annotations for visual object detection. In *AAAI Workshops on Human Computation*, 2012.
- [28] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deep-Face: Closing the gap to human-level performance in face verification. In *Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [29] A. Vedaldi and A. Zisserman. Efficient additive kernels via explicit feature maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, 34(3):480–492, 2012.
- [30] A. Vezhnevets and V. Ferrari. Associative embeddings for large-scale knowledge transfer with self-assessment. In *Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [31] F. Wilcoxon. Individual comparisons by ranking methods. *Biometrics bulletin*, 1945.
- [32] C. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In *Conference on Neural Information Processing Systems (NIPS)*, 2001.
- [33] K. Zhang, I. W. Tsang, and J. T. Kwok. Improved Nyström low-rank approximation and error analysis. In *International Conference on Machine Learning (ICML)*, 2008.